Deloitte.

Together makes progress

Engineering the agentic enterprise

How software, AI, and cybersecurity are converging to redefine enterprise agility and resilience

Enterprise software is at an inflection point

Enterprise software is entering a significant new era. Many organizations are moving beyond traditional, standard packaged applications and configurable software-as-a-service (SaaS) toward "SaaS+" models, which are platforms augmented by agent-enabled customization. The next horizon most certainly will be agentic architectures, where autonomous software agents can coordinate workflows, make decisions, and act with human guidance and oversight.

This evolution is not an "either/or" proposition between packaged and custom, or human-led and agent-led. It is an "and" approach that can help enterprises determine which models are right for their needs. However, the challenge for executives is to understand how best to integrate these models so they strike an effective balance between innovation, security, and governance.

The shift: from packaged to SaaS+ to agentic only

It's crucial to appreciate how the shift unfolded to understand where it's going. Historically, organizations relied on packaged or custom, off-the-shelf (COTS) software, customizing it as little as possible to avoid unnecessary complexity. They also historically avoided customizing the solutions to accommodate their unique processes due to the restrictive nature or cost of customization. The rise of SaaS enabled faster deployment, regular updates, and configurable workflows while keeping complexity in check. Today, there is a shift toward SaaS+, which inserts agentic capabilities into standard platforms, thereby embedding intelligence, automation, and customization within core systems.

The next stage, and what we believe is the future of software engineering, is agentic-only, which is the process of building systems with autonomous agents that can analyze, act, and collaborate across environments. With agentic artificial intelligence (AI), agents can address the unique processes of individual organizations to meet their business needs. This abstraction over COTS solutions gives organizations the flexibility they need. For example, agents can autonomously interpret business goals, invoke tools and processes, and execute workflows across systems, often with more efficiency than traditional automation. Agents can also provide persona-based "digital twins" to execute a variety of business functions that will not require corresponding growth in headcount.

What's driving the shift?

A confluence of technological, economic, and strategic forces is driving the shift. The cost of software development and customization has decreased significantly because of advances in model orchestration, agentification and co-piloting vis-a-vis the development lifecycle, low-code frameworks, and agent architectures. Simultaneously, foundation models and multi-agent systems have advanced, which leads to high-level performance at potentially lower cost. The calculus of build-or-buy decisions is changing as a result. Now, it's possible to create flexible, tailored solutions in-house without having to depend on software vendors.

Of course, the pace at which enterprises undergo the shift—or even whether they do or not—will depend on each organization's risk tolerance and regulatory needs, as well as their data governance and maturity state. For instance, some processes may remain on standard SaaS due to stability concerns and because of processes and systems of record that still require deterministic rule sets and solutions (e.g., core financial systems). Others—especially data-intensive and rules-based processes and complex workflows with dynamic dependencies—are a good fit for agentic augmentation. The most successful organizations are able to pragmatically integrate all three approaches and deploy each, or a combination, where they can deliver the most business value rather than simply replacing the old with the new.

Key enablers: Infrastructure, observability, and security

For organizations to operate effectively across a continuum of packaged, SaaS+, and agentic AI, three enablers are crucial: modern infrastructure, advanced observability, and robust cybersecurity.

Infrastructure

A cloud-native, edge, and automated infrastructure is the foundation for the agility that enables choice. Agentic applications require dynamic compute capacity, near real-time data access, and flexible scalability. Techniques like containerization, orchestration, and infrastructure-as-code approaches also allow organizations to efficiently deploy, scale, and integrate agents across hybrid environments. Where latency or local processing capabilities matter, edge computing becomes essential. Without this agility, agentic systems could remain proofs of concept rather than valuable, scalable realities.

Observability

Observability enables organizations to monitor how agents operate, make decisions, and impact broader systems. However, observability is more than simple monitoring. It can provide insight into agent behavior, workflow efficiency, and operational anomalies to enable teams to spot issues and remediate them before they become critical. Further, metrics on decision quality, tool use, and escalation patterns can enable continuous improvement and reveal areas that could benefit from human oversight. Finally, observability fosters trust in systems that are partially autonomous "black boxes," which is a nonnegotiable condition for scaling agentic operations to the enterprise.

Security

Because agentic architectures redefine the security environment, modern approaches such as zero trust, identity-aware networking, and application programming interface (API)-level protection are mandatory. Fundamentally, each autonomous agent represents a new identity within the enterprise environment that requires authentication, authorization, and auditing akin to that of a human user. Complexity is increased with autonomous agents working in a multi-agent system to achieve a business goal. Agents can spawn ephemeral agents and control, auditing, discovery, and life cycle management of child agents.

Protecting these environments involves several critical steps:

- **Zero trust.** Each interaction, whether human or agent initiated, must be verified. Techniques like least-privilege access and micro-segmentation can reduce risk.
- **Agent and API governance.** Just like humans, each agent must adhere to defined access parameters, credentials, and life cycle management policies such as registration, auditing, and retirement.
- **Data security and governance.** Agents often require access to large, disparate volumes of enterprise data. So, it's essential to practice effective data handling, lineage tracking, and encryption.
- **Life cycle control.** Agents learn and iterate as they evolve. Their decisions and outputs must be completely traceable, explainable, and constantly subject to review.

These factors are fundamental to an enterprise's ability to innovate securely. Infrastructure enables performance, observability fosters trust, and security maintains control. Taken together, they can determine whether agentic AI introduces more vulnerability or delivers sustainable value.

Engineering and cyber must collaborate to build and secure the future

The success of agentic AI depends immensely on the creation of a trusting, collaborative partnership between software engineering and cybersecurity. Both may need to shed their histories as independent, sometimes rival, disciplines to become partners that are jointly responsible for design, delivery, and safeguarding intelligent enterprise systems. Trust is the cornerstone of successful AI adoption.

Collaboration to align efficiency with security

Many enterprise functions are ripe for agentic augmentation, but it's not a panacea. Operational, infrastructure, and customer-facing activities are most likely to be suitable for agentic AI.

Back-office processes, IT operations, infrastructure management, software testing, and workflow-intense environments are strong candidates. Agents can generate and validate code, manage configurations, triage incidents, automate testing, or deploy approval chains, which can streamline operations while maintaining compliance.

Likewise, in customer-facing areas, agents can coordinate marketing campaigns, manage multichannel interactions, or personalize product recommendations. Within IT, they can support DevSecOps pipelines, automate deployment checks, and continuously monitor compliance.

The need for collaboration between engineering and cybersecurity is clear. Engineering and cybersecurity can work together to determine where agentic AI can offer the most return on investment. Engineering can identify tasks for potential automation, while cybersecurity can develop mechanisms to ensure that solutions maintain compliance and manage risk effectively. They can also prioritize ways in which efficiency and security can be jointly improved.

The role of advanced engineering

Modern engineering practices make agentic development possible at scale. Modular, API-first design; event-driven systems; and continuous integration/continuous delivery pipelines all support safe experimentation and quick delivery of applications. Further, teams can build reusable frameworks, implement model and logic version control, and detect incorrect configurations or anomalous behavior. Crucially, human software engineers remain central to orchestrating these complex systems by providing oversight, ensuring adaptability, and maintaining alignment with ethical standards at every stage. Advanced testing environments and sandboxes are also essential to verify that agents act within appropriate, explicit boundaries. The drive for engineering excellence can help ensure that autonomy enhances, rather than undercuts, reliability.

The role of advanced cybersecurity

Cybersecurity practices are quickly having to adapt to a world where machines act on behalf of humans. As with humans, each agent needs a certifiable identity, defined privileges, and monitoring to ensure that tasks are completed correctly. Therefore, cyber teams should implement continuous authentication, effective credentialing, and threat detection processes that are customized to expected agent behavior. Cross-domain (intercompany) agent orchestration creates more complexity as agents need to be aware of agent impersonation. A trusted registry to discovery and consume agents becomes a core part of the trust infrastructure.

Cyber should also develop escalation protocols for those times that agents encounter unclear decision options or potential security threats. Visibility into data flows between agents, APIs, and models is also critical to prevent security breaches or misuse. Finally design techniques like isolation, encryption, and immutable logging must extend into all stages of the agent life cycle—from creation to retirement.

The goal: Engineering and cyber collaboration

The collaboration between engineering and cyber should be intentional. They must work together to embed an ethos of governance, scalability, and security into the organization's efforts to implement agentic AI. Engineering should involve cybersecurity from the earliest design phases so that appropriate risk considerations can inform decisions and choices from architecture through deployment.

Conversely, it's essential that cyber teams understand engineering pipelines, observability methods and dashboards, and development schedules to ensure security without hindering progress. Joint playbooks, shared information, and collaborative incident response can help close the loop between development and security.

Fundamentally, the alignment of engineering and cyber to form a collaborative partnership can transform security from a development constraint to an accelerator. When engineering and cyber collaborate, organizations can deploy agentic solutions confidently because they understand that innovation and protection advance together. The result is a secure foundation for intelligent automation that can scale as needed to provide outcomes that meet business goals.

The partnership payoff

Organizations that pursue a balanced adoption strategy across packaged, SaaS+, and agentic architectures can anticipate gains:

- **Readiness for the future.** Enterprises equipped with agentic capabilities can adapt more rapidly as technologies and markets evolve.
- **Greater agility and speed to market.** Agents can automate multi-step workflows, shorten release cycles, and accelerate delivery.
- **Faster innovation.** By automating repetitive development and operational tasks, teams gain more capacity to innovate.
- **Operational resilience.** Agentic systems can operate continuously, self-recover from disruptions, and improve reliability.
- **Enhanced customer experience.** Streamlined internal operations lead to faster, more accurate, and personalized customer interactions.
- **Cost efficiency and resource optimization.** Automation reduces manual effort, improves utilization, and maximizes return on technology investments.

These benefits are not theoretical; they are measurable. Organizations that approach the transition as an orchestrated journey, rather than a wholesale transformation, will likely see the greatest return.

Deloitte's approach to helping organizations navigate the shift

Deloitte can help clients negotiate this transition with a structured approach that integrates engineering excellence, cybersecurity discipline, and strategic foresight. We begin by assessing organizational readiness, identifying high-value use cases, and defining a roadmap that's based on your unique environment and needs.

Our teams combine deep industry knowledge with cross-domain experience in cloud engineering, agentic architectures, and enterprise security. We design and build secure, customized solutions that align technology strategy with business outcomes. Our approach is collaborative: We work alongside client teams to accelerate adoption, reinforce governance, and build internal capability.

Deloitte's strength lies in the breadth of knowledge and services we bring. We link advisory, engineering, and cyber practices to help organizations deploy agentic systems that are secure, scalable, and aligned with enterprise goals. Our focus remains constant: delivering measurable business value through secure innovation.

Progress is a process of managed, secure evolution

The emergence of agentic AI heralds an evolution in how enterprises design, deploy, and manage software. Success in this new era will require thoughtful planning—deciding where to maintain packaged solutions, where to move to SaaS+, and where to implement fully agentic architectures—and whether and when to deploy all three.

Leaders must act intentionally. They should determine their readiness to meet their goals, align engineering and cyber functions into a collaborative team, and establish the governance necessary for successful, practical deployment. Organizations that don't embrace the challenges of the agentic revolution can risk falling behind more agile competitors and increasing their exposure to operational and security risks.

By approaching agentic adoption as a process that's grounded in engineering discipline, cybersecurity consistency, and strategic vision, leaders can turn potential into better performance and gain a competitive edge in the era of the intelligent enterprise.

Continued reading

Engineering a culture of technology resilience: How enterprises can proactively thrive amid uncertainty.

Authors

Faruk Muratovic
Principal
AI & Engineering Strategy & Services Leader
Deloitte Consulting LLP
famuratovic@deloitte.com

Vikram Kunchala
Principal
Digital Trust & Privacy Leader
Deloitte & Touche LLP
vkunchala@deloitte.com

As used in this document, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.